

CCPA Patent Disclosure Form

This form is provided by CCPA in an effort to obtain enabling written disclosure for the purpose of preparing patent cases on behalf of our clients.

Please take the time to fill it out properly as the provision of clear and concise disclosure will speed the process of preparing and filing your case. In addition a signed disclosure document filed with the patent office can effectively remove references anticipating our invention in the prosecution of the case thereby enhancing our chances of obtaining patents. In most cases, verbal disclosure or short e-mail messages are inadequate forms of disclosure and should be avoided.

To fill out the form correctly, follow each set of instructions provided with each heading.

Title of Invention

This section is simply a brief descriptive title of the invention.

Insert title here: Context selection mechanism for pre-loading packet information in the background

Inventors

We will need the residence address, mailing address, full legal names and citizenship of each inventor at the time of submission of the disclosure.

Enrique Musoll

Related inventions known or authored by you or your company

This section should list any prior patents known to you or patents that you have already filed if the <u>present invention</u> depends on them for successful practice.

- [1] PA3818
- [2] "Design Tradeoffs for Embedded Network Processors", by Tilman Wolf, Mark A. Franklin and Edward W. Spitznagel. WUCS-00-24, July 20, 2000. Washington University in Saint Louis.

Background

This section is used to describe "the state of the art" before being improved or enhanced with your invention. It should include a brief summarization of existing technologies if any that the present invention improves upon or replaces, a description of any specific problems with "the way the art is practiced now", and a very brief statement of what is needed to improve or replace the existing art. Include references by U.S. patent number any closely related patents discovered during any prior-art searches

Begin Background here:

This patent disclosure presents the algorithm implemented in XCaliber that chooses an available context from the pool of 8 contexts whenever a packet requires automatic activation or whenever a context is requested by software.

The algorithm relies on the clustering partition of functional units in the processor core and selects the context that is not being used by the processor that maximizes the use of the functional units in both the clusters.

This algorithm was first mentioned in [1] and this patent disclosure refines it and provides other improvements that could be implemented in a future generation.

This disclosure also presents improvements on the algorithm implemented in XCaliber.

Description of Invention

This section should explain the basic apparatus and method of practicing your invention according to a preferred state. If certain apparatus of the invention is not known in the prior art then indicate so. If a method of the present invention is not known in prior art then indicate so. If certain methods and apparatus are known in prior art then they do not have to be greatly detailed. However any new subject matter novel over the prior art should be fully explained and represented by drawings and/or sketches.

Begin description here:

Introduction

There exists a hardware in the Packet Management Unit (PMU) of XCaliber that:

- a) pre-loads an available context with an available context with information of packets that require automatic activation
- b) provides an available context to the processing core

This hardware is henceforth named Register Transfer Unit or RTU.

In either case of the two above cases, it is beneficial that a context selected by this hardware so that it maximizes the performance of the processing core (henceforth named SPU). In other words, it is important that the context is selected so that the different streams running on the processing core conflict the least in their attempt to use the different shared functional units.

The context selection algorithm implemented in XCaliber relies on the clustering of functional units of the processing core to choose an available context that maximizes the use of the functional units.

Context States

A context in XCaliber can be in one of two states: PMU-owned or SPU-owned. The ownership of a context changes when the current owner releases the context. The PMU releases a context to the SPU whenever:

- 1) The RTU has finished pre-loading the information of a packet into the context
- 2) The SPU requests a context to the RTU
- 3) All 8 context are PMU-owned.

This patent disclosure covers the algorithm that the RTU implements to select a context in these cases.

The SPU releases a context to the RTU when the SPU executes the RELEASE instruction, an XStream proprietary instruction.

Context Selection Algorithm

There are eight functional units in the SPU core. However, a stream can only issue instructions to a fixed set of four functional units. The stream running on context 0-3, can

only issue instructions to the functional units located in cluster 0, whereas an stream running on context 4-7 can only issue instruction to the functional units located in cluster 1.

The RTU may own several contexts at a given time. Logic is required to select one of these contexts when a pre-load is performed, or when a context must be provided to the SPU. The goal of the logic is to balance the pressure in the functional units, i.e. to spread the requests for functional units evenly across both clusters.

The selection logic implemented in XCaliber has as input the state of the different contexts. The following table of numbers specifies the truth table of the logic. Each number is associated to a possible combination of SPU/PMU-owned context. For example, the first number corresponds to the combination '00000001', meaning that context number 0 is PMU owned and context numbers 1 to 7 are SPU owned. The second number corresponds to combination '00000010', the third to combination '00000011', and so forth up to combination '11111110' (note that combinations '00000000' and '11111111' are not applicable. The first one implies that there is no PMU-owned context and, therefore, no selection will be performed. The second one implies that all the contexts are PMU owned, and this will never occur since at least one context will always remain SPU owned).

For example, the 19^{th} combination ('00010011') has associated number 3 (or '00000011') in the previous list, which means that context 0 and 1 can be selected by the

select logic. Context 4 could also be selected (since it is PMU owned), but it is not be best choice to balance the use of the functional units in the SPU.

Improvements to the algorithm

There are several improvements to the algorithm described above. In [2] an improvement is hinted, but for a hardware-based context selection for chip multiprocessing architectures, not multi-streaming architectures. In [2] the authors propose that the context to be selected is determined based on the type of the packet that the selected context is going to be process. The goal is then to select a context that has previously processed a packet of the same type as the new packet, so that the instruction and data locality is improved (the old instructions and data values might be reused by the processing of the new packet).

The algorithm described in this patent disclosure and implemented in XCaliber can be improved as follows:

- a) By taking into account the stall status information of the streams. If 3 streams within a cluster are all stalled waiting for a value in external memory, the available context is a perfect candidate for selection since it will be able to issue instructions to all the functional units in its cluster.
- b) By predicting how much time the processing of the corresponding thread will take. If this prediction is accurate, the RTU can perform a better context selection by forcing a mix of long and short streams into the same cluster. For example, if
 - a. context 0 (cluster 0) is running a short stream, and
 - b. context 4 (cluster 1) is running a long one, and
 - c. the rest of the contexts are PMU-owned

the RTU will select one of the contexts in cluster 0 if the packet is to be processed by a long stream, whereas it will select a context in cluster 1 of the stream is short.

c) By predicting the distribution of instruction types that the processing of the packet will have. If the clusters in the processing core had asymmetrical

composition of functional units, the RTU can select a context in the cluster that has the most appropriate type of functional units. For example, if the stream is going to execute a lot of multiplications and the single multiplier unit is in cluster 0, then the RTU is forced to pick up a context in cluster 0. If there is a fast (but costly in area) multiplier in cluster 1 and a slow one in cluster 0, and no stream is expected to be using any of the multipliers, then RTU will select a context within cluster 1.

Please have all inventors sign the disclosure and mail a hard copy to CCPA for participation in the document disclosure program. Also e-mail to Mark Boys markboys@centralcoastpatent.com and CC Don Boys rexboys@centralcoastpatent.com Your cooperation in the filling and return of this form will expedite the processing of your application and increase our chances of obtaining a patent for your invention. Mark Boys, CCPA